Network calculus for parallel processing

G. Kesidis, B. Urgaonkar, Y. Shan, S. Kamarava
CSE and EE Depts
The Pennsylvania State University
{gik2,buu1,yxs182,szk234}@psu.edu

J. Liebeherr
ECE Dept
University of Toronto
jorg@comm.utoronto.ca

arXiv:1409.0820v2 [cs.PF] 31 Jan 2015

I. Introduction

Multi-stage parallel data processing systems are ubiquitous in cloud computing environments. In a single parallel processing stage, a job is partitioned into tasks (*i.e.*, the job is "forked" or the tasks are demultiplexed); the tasks are then worked upon in parallel. Within parallel processing systems, there are often processing "barriers" (points of synchronization or "joins") wherein all component tasks of a job need to be completed before the next stage of processing of the job can commence. The terminus of the entire parallel processing system is typically a barrier. Thus, the latency of a stage (between barriers or between the exogenous job arrival point to the first barrier) is the greatest latency among the processing paths through it (*i.e.*, among the tasks comprising that stage). Google's multi-stage MapReduce [6] (especially its open-source implementation Apache Hadoop [15]) is a very popular such framework. However, numerous other systems exhibit similar programming patterns and our work is relevant to them as well [16], [13], [8].

In MapReduce, jobs arrive and are partitioned into tasks. Each task is then assigned to a mapper for initial processing. The results of mappers are transmitted (shuffled) to reducers. Reducers combine the mapper results they have received and perform additional processing (a final stage after the reducers may simply combine their results). The workloads of the reducer tasks may be *unrelated* to those of their "tributary" mapper tasks. A barrier exists before each reducer (after its mapper-shuffler stage) and after all the reducers (after the reducer stage).

To achieve good interleaving of the principal resources consumed by the mapper (CPU/memory) and the shuffler (network bandwidth), these stages are made to work in a pipelined manner wherein the shuffler transmits partial results created by the mapper (as they are generated) rather than waiting for a mapper to entirely finish its task. Of course, the shuffler must "follow" the mapper at all times

in the sense of being able to send only what results the mapper has generated so far [11]. On the other hand, the barrier between the shuffler stage and the reducer stage is a strict one - a reducer may not begin any processing until all of the shuffler stage's work is done.

Our goal is to develop a performance model for such applications. As a first step toward this, we consider a single parallel processing stage. Our approach can be extended to create a model with separate queues for the mapper, the shuffler, and the reducer stages (or even more stages), but we restrict our attention to the interaction between the shuffler and the reducer stages. Specifically, each processor/server (and associated job queue) in our model represents a mapper stage.

In this paper, we focus on the Mapper stage, where an initial job scheduling would be in play to achieve a "bounded burstiness" of the aggregate workload. We prove two claims using "network calculus" for parallel processing systems. We then numerically evaluate the "generalized" (strong) stochastic burstiness bound (gSBB) of a publicly disseminated workload trace of a Facebook data-center.

II. Related Work

There is substantial prior work on fork-join queueing systems particularly involving underlying Markov chains, *e.g.*, [12]. Our primary result is significantly simpler and somewhat more general than those based on Markovian models of workload. Of course, there is also an enormous literature on parallel processing systems in general. Typically, parallel processing systems employ robust load balancing to minimize synchronization delays at the barriers. To this end, load balancing could proactively estimate throughputs along the parallel processing paths and proportionately size the workloads from tasks fed to them. As an example of the many proposed reactive/dynamic mechanisms, "straggler" (deemed excessively delayed) tasks at barriers can be restarted or the entire job can be interrupted and restarted or additional can be allocated (*e.g.*, more parallelism). See [10], [7], [2], [1] for recent discussions on the online management specifically for a MapReduce parallel processing system. Some recent work on MapReduce systems [4], [11] focuses on the the pipelining between the mapper and shuffler, the latter formulating a proactive *scheduling* problem that jointly considers individual job workloads of both mapper and shuffler (assuming the shuffler load is known *a priori*).

III. Single-stage, fork-join system

Consider single-stage fork-join (parallel processing) system, modeled as a bank of K parallel queues, with queue-k provisioned with service/processing capacity s_k. Let A be the cumulative

input process of work that is divided among queues so that the kth queue has arrivals a_k and departures d_k in such a way that $\forall t \geq 0$,

$$A(t) \;=\; \sum_k a_k(t).$$

Define the virtual delay processes for hypothetical departures from queue k at time $t \geq 0$ as

$$\delta_k(t) \;=\; t - a_k^{-1}(d_k(t)),$$

where we define inverses a_k^{-1} of non-decreasing functions a_k as continuous from the left so that $a_k(a_k^{-1}(v)) \equiv a_k^{-1}(a_k(v)) \equiv v$.

In following definition of the cumulative departures, D, the output is determined by the most lagging (straggling) queue/processor: $\forall t \geq 0$,

$$\begin{aligned} D(t) &= A(t - \max_k\{\delta_k(t)\}) \\ &= A\left(\min_k\{a_k^{-1}(d_k(t))\}\right). \end{aligned}$$

Note that in the case of continuous, fluid arrivals (*e.g.*, piecewise linear A), this definition of departures D corresponds to periods of continual, possibly perpetual, barriers (synchronization times). In the case of discrete arrivals (piecewise constant A with jump discontinuities at arrival instances), then the barriers are discrete.

A queue q with service s has a (non-negative, non-decreasing) at least a service-curve $s_{\min}$ (*i.e.*, $s \gg s_{\min}$) if for all cumulative arrivals a and all time t its cumulative departures

$$d(t) \;\geq\; (s_{\min} * a)(t) \;:=\; \min_{v \leq t}\{a(v) + s_{\min}(t-v)\}.$$

Define the convolution ($*$) identity as

$$u_\infty(t) \;=\; \begin{cases} 0 & \text{if } t \leq 0 \\ +\infty & \text{if } t > 0 \end{cases}$$

and

$$\begin{aligned} d_{\max,k} \;=\; &\min\{z \geq 0 \;:\; \forall x \geq 0, \\ &s_{\min,k}(x) \geq (b_{\text{in},k} * \Delta_z u_\infty)(x) = b_{\text{in},k}(t-z)\} \end{aligned} \tag{1}$$

where:

- the delay operator $(\Delta_d g)(t) \equiv g(t-d)$,
- $a_k \ll b_{\text{in},k}$, *i.e.*, a_k conforms to the non-negative, non-decreasing burstiness curve (traffic envelope) $b_{\text{in},k}$,

$$\forall t,\ a_k(t) \ \leq \ (a_k * b_{\text{in},k})(t),$$

- $d_{\max,k}$ is the largest horizontal difference between $b_{\text{in},k}$ and $s_{\min,k}$ [5].

Claim 1. *If the kth queue has at least a service curve of $s_{\min,k}$ and arrivals $a_k \ll b_{\text{in},k}$, then for all $t \geq 0$,*

$$D(t) \ \geq \ A(t - \max_k\{d_{\max,k}\}).$$

Remark: This claim simply states that the maximum delay of the whole system (from A to D) is the maximum delay among the queues. Equivalently, the service curve from A to D is at least $\Delta_d u_\infty$, where $d := \max_k d_{\max,k}$.

Proof: By Equation (1), $\forall t \geq v \geq 0$ and $\forall k$,

$$\begin{aligned} s_{\min,k}(t-v) &\geq b_{\text{in},k}(t-v-d_{\max,k}) \\ &\geq a_k(t-d_{\max,k}) - a_k(v) \end{aligned}$$

Thus, $\forall t \geq v \geq 0$ and $\forall k$,

$$\begin{aligned} a_k(v) + s_{\min,k}(t-v) &\geq a_k(t-d_{\max,k}) \\ \Rightarrow \quad (a_k * s_{\min,k})(t) &\geq a_k(t-d_{\max,k}) \\ \Rightarrow \quad a_k^{-1}((a_k * s_{\min,k})(t)) &\geq t - d_{\max,k} \end{aligned}$$

where we have used the fact that, $\forall k$, a_k are nondecreasing. Thus,

$$\begin{aligned} D(t) &= A\left(\min_k\{a_k^{-1}(d_k(t))\}\right) \\ &\geq A\left(\min_k\{a_k^{-1}((a_k * s_{\min,k})(t))\}\right) \\ &\geq A\left(\min_k\{t - d_{\max,k}\}\right) \\ &= (A * \Delta_d u_\infty)(t), \end{aligned}$$

where we have used the fact that A is nondecreasing. □

We now consider a stationary stochastic model of this single-stage system. To simplify matters, we assume the workload process A has strong ("generalized") stochastically bounded burstiness (gSBB) [9], and leave to future work generalizations to non-stationary settings assuming only (weak) stochastically bounded burstiness [14] or bounded log moment-generating function [3].

Claim 2. *In the stationary regime at time $t \geq 0$, if*

A1 *service to queue k, $s_k \gg s_{\min,k}$ where*

$$\forall v \geq 0, \quad s_{\min,k}(v) \;:=\; v\mu_k;$$

A2 *$\forall k$, $\exists$ small $\varepsilon_k > 0$ such that $\forall v \leq t$*

$$\left|a_k(t) - a_k(v) - \frac{\mu_k}{M}(A(t) - A(v))\right| \;\leq\; \varepsilon_k \quad \textit{a.s.},$$

where $M := \sum_k \mu_k$;

A3 *the total arrivals have strong stochastically bounded burstiness [9],*

$$\mathsf{P}(\max_{v \leq t}\{A(t) - A(v) - M(t-v)\} \geq x) \;\leq\; \Phi(x),$$

where Φ decreases in $x > 0$;

then $\forall x > 2M \max_k\{\varepsilon_k/\mu_k\}$,

$$\mathsf{P}(A(t) - D(t) \geq x) \;\leq\; \Phi(x - 2M \max_k\{\varepsilon_k/\mu_k\}).$$

Remark: By A2, the mapper divides arriving work roughly proportional to minimum allocated service resources μ_k to queue k, *i.e.*, strong load balancing.

Proof:

$$\begin{aligned}
&\mathsf{P}(A(t) - D(t) \geq x) \\
&= \mathsf{P}(A(t) - A(\min_k\{a_k^{-1}(d_k(t))\}) \geq x) \\
&= \mathsf{P}(\min_k\{a_k^{-1}(d_k(t))\} \leq A^{-1}(A(t) - x) =: t - z) \\
&= \mathsf{P}(\exists k \text{ s.t. } d_k(t) \leq a_k(t - z)) \\
&= \mathsf{P}(\exists k \text{ s.t. } a_k(t) - d_k(t) \geq a_k(t) - a_k(t - z) =: x_k), \\
&\leq \mathsf{P}(\exists k \text{ s.t. } \max_{v \leq t}\{a_k(t) - a_k(v) - (t - v)\mu_k\} \geq x_k)
\end{aligned}$$

where we have used the fact that A and the a_k are nondecreasing (cumulative arrivals) and the inequality is by A1. Also, we have defined non-negative random variables z and x_k such that $\sum_k x_k = x = A(t) - A(t - z)$. So by using A2 then A3, we get

$$\begin{aligned}
&\mathsf{P}(A(t) - D(t) \geq x) \\
&\leq \mathsf{P}(\exists k \text{ s.t. } \max_{v \leq t}\{\frac{\mu_k}{M}(A(t) - A(v)) + \varepsilon_k - (t - v)\mu_k\} \\
&\qquad \geq \frac{\mu_k}{M}x - \varepsilon_k) \\
&= \mathsf{P}(\exists k \text{ s.t. } \max_{v \leq t}\{(A(t) - A(v)) - (t - v)M\} \\
&\qquad \geq x - 2\frac{M}{\mu_k}\varepsilon_k) \\
&= \mathsf{P}(\max_{v \leq t}\{(A(t) - A(v)) - (t - v)M\} \\
&\qquad \geq x - 2M \max_k\{\varepsilon_k/\mu_k\}) \\
&\leq \Phi(x - 2M \max_k\{\frac{\varepsilon_k}{\mu_k}\}).
\end{aligned}$$

□

ACKNOWLEDGEMENTS

This work was supported in part by a Cisco Systems URP gift and NSF CAREER grant 0953541.

REFERENCES

[1] G. Ananthanarayanan, A. Ghodsi, S. Shenker, and I. Stoica. Effective straggler mitigation: Attack of the clones. In *Proc. USENIX NSDI*, Lombard, IL, Apr. 2013.

[2] G. Ananthanarayanan, S. Kandula, A.G. Greenberg, I. Stoica, Y. Lu, B. Saha, and E. Harris. Reining in the Outliers in Map-Reduce Clusters using Mantri. In *Proc. USENIX OSDI*, Vancouver, Oct. 2010.

[3] C.-S. Chang. Stability, queue length, and delay of deterministic and stochastic queueing networks. *IEEE Trans. Automatic Control*, 39:913931, May 1994.

[4] F. Chen, M. Kodialam, and T.V. Lakshman. Joint Scheduling of Processing and Shuffle Phases in MapReduce Systems. In *Proc. IEEE INFOCOM*, 2012.

[5] Rene L. Cruz. SCED+: Efficient Management of Quality of Service Guarantees. In *Proc. IEEE INFOCOM*, pages 625–634, 1998.

[6] J. Dean and S. Ghemawat. MapReduce: Simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, 2008.

[7] M. Ead, H. Herodotou, A. Aboulnaga, and S. Babu. PStorM: Profile storage and matching for feedback-based tuning of MapReduce jobs. In *Proc. Conf. on Extending Database Technology*, March 2014.

[8] Google Cloud Platform Blog, Sneak peek: Google Cloud Dataflow, a Cloud-native data processing service. http://googlecloudplatform.blogspot.com/2014/06/sneak-peek-google-cloud-dataflow-a-cloud-native-data-processing-service.html, June 26, 2014.

[9] Y. Jiang, Q. Yin, Y. Liu, and S. Jiang. Fundamental calculus on generalized stochastically bounded bursty traffic for communication networks. *Computer Networks*, 53(12):2011–2021, 2009.

[10] M. Li, L. Zeng, S. Meng, J. Tan, L. Zhang, N.C. Fuller, and A.R. Butt. Online: MapReduce Online Performance Tuning. In *Proc. HPDC*, 2014.

[11] M. Lin, L. Zhang, A. Wierman, and J. Tan. Joint optimization of overlapping phases in MapReduce. *Perform. Eval.*, 70(10):720–735, 2013.

[12] R. Nelson and A.N. Tantawi. Approximate analysis of fork/join synchronization in parallel queues. *IEEE Trans. Computers*, 37, 1988.

[13] C. Reiss, A. Tumanov, G.R. Ganger, and R.H. Katz. Heterogeneity and Dynamicity of CLouds at Scale: Google Trace Analysis. In *Proc. ACM SOCC*, Oct. 2012.

[14] D. Starobinski and M. Sidi. Stochastically bounded burstiness for communication networks. In *Proc. IEEE INFOCOM*, March 1999.

[15] T. White. *Hadoop: The Definitive Guide*. O'Reilly Media, Inc., 1st edition, 2009.

[16] Y. Yu, M. Isard, D. Fetterly, M. Budiu, U. Erlingsson, P.K. Gunda, and J. Currey. DryadLINQ: A System for General-Purpose Distributed Data-Parallel Computing Using a High-Level Language. In *OSDI*, pages 1–14, 2008.